AFOSR·TR· 9 0 - 0 2 8 8

AD-A219 230

THE PERFORMANCE OF WAVELETS FOR DATA COMPRESSION IN SELECTED MILITARY APPLICATIONS

FINAL REPORT

prepared by
Aware, Inc.
124 Mt. Auburn Street, Suite 310
Cambridge, MA 02138
(617) 354-3311

with the assistance of
Atlantic Aerospace Electronics Corporation
6404 Ivy Lane, Suite 300
Greenbelt, MD 20770
(301) 982-5275

Aware Technical Report AD 900223
23 February 1990

90 03 13 072

REPORT DOCUMENTATION PAGE

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	23 February 1990	Final Report, 28 Aug 89 to 27 Dec 89

4. TITLE AND SUBTITLE

THE PERFORMANCE OF WAVELETS FOR DATA COMPRESSIONS IN SELECTED MILITARY APPLICATIONS

5. FUNDING NUMBERS

F49620-89-C-0122
61101E   7145/00

6. AUTHOR(S)

Howard L. Resnikoff

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Aware, Inc.
124 Mt. Auburn Street, Suite 310
Cambridge, MA   02138

8. PERFORMING ORGANIZATION REPORT NUMBER

AFOSR·TR· 00-0288

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFOSR/NM, Bldg, 410
Bolling AFB, DC   20332-6448

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

AFOSR-
F49620-89-C-0122

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Wavelets provide a new mathematical and computational approach to representing image data. A wavelet basis is a complete orthonormal system of functions in terms of which image data can be represented. The low computational complexity of the wavelet transform and the bounded support of the wavelet basis functions off high retention of information that typically corresponds to image features, and low workload for computing, compressing, and reconstituting image data. Aware conducted computational experiments to determine the contribution of wavelets to solving two classes of practical problem: 1. Position location by matching observations to a stored image. 2. Identification of objects of specific size or characteristics and sensory clutter. The results suggest that compression using wavelets preserves more than enough information to register position correctly using position reconstituted reference images that have been compressed in excess of 100:1. A variety of analysis made to develop an indication of the relationship between information loss and compression are reported.

14. SUBJECT TERMS	15. NUMBER OF PAGES
	56
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSSIFIED	UNCLASSIFIED	

Executive Summary

This is the final report on research to "demonstrate performance of wavelets for data compression in selected military applications." This work was sponsored by the Defense Advanced Research Projects Agency under DARPA order number 7125 and monitored by the Air Force Office of Scientific Research under contract number F49620-89-C-0122.

Within the constraints of the project, the experiments showed the ability to correctly determine location for navigational purposes using reference image "maps" reconstituted from compressed image files for compression ratios between 15:1 and 117:1. Accurate location using images reconstituted from compressed files was achieved in all cases at compressions of 117:1. Preliminary demonstrations with noise and scaling properties showed that the ability to correctly locate was retained at the same compressions in the presence of more than 20% noise or a scale variation of up to 20%.

These results appear to be consequences of the excellent preservation of transients by Aware's wavelet compression methods which resulted in a highly predictable slow decline in the signal-to-noise ratio of the Laplacian of the correlation function. As these methods are known to involve fewer mathematical operations than traditional compression methods and to be both parallellizable and pipelineable, they require less computer power than more established methods. they can be easily calculated in real time, and they will fit easily onto special purpose chips, if desired.

Based on the achievements thus far, wavelet compression could provide an order of magnitude improvement in compression of reference images for navigational purposes, capable of being fielded very rapidly at low cost. Specific follow-on investigations are proposed

In addition, some experiments were conducted using wavelet techniques to find objects in clutter. While the methods here are far earlier in the development process, they indicated some promising directions for future research.

PERFORMANCE OF WAVELETS FOR DATA COMPRESSION IN SELECTED MILITARY APPLICATIONS

Contents

Volume I: Final Report

Volume II: Supplementary Tables and Graphs

Introduction and Overview

This is the final report on research to "demonstrate performance of wavelets for data compression in selected military applications." This work was sponsored by the Defense Advanced Research Projects Agency under DARPA order number 7125 and monitored by the Air Force Office of Scientific Research under contract number F49620-89-C-0122.

This report is in two volumes. Volume I is the report. Volume II provides supplementary tabular and graphic exhibits.

Wavelets provide a new mathematical and computational approach to representing image data. A wavelet basis is a complete orthonormal system of functions in terms of which image data can be represented. The low computational complexity of the wavelet transform and the bounded support of the wavelet basis functions offer high retention of information that typically corresponds to image features, and low workload for computing compressing and reconstituting image data.

Aware conducted computational experiments to determine the contribution of wavelets to solving two classes of practical problem:

(1) position location by matching observations to a stored image and

(2) identification of objects of specific size or characteristics amid sensory clutter.

The images employed both sets of experiments were (512 x 512) pel x 8 bits per pel, gray scale, digitized images of aerial photographs and simulated radar output. Atlantic Aerospace Electronics Corporation selected the images as realistic examples, supplied them to Aware, suggested performance measures for assessing the results of the experiments, and participated in the preparation of this final report. Aware, Inc. carried out all other tasks.

In the position location experiments, images of terrain (aerial photographs) were compressed using algorithms based upon the wavelet transform that had been previously developed by Aware. The compressed images were then reconstituted and tests were run to determine the accuracy with which a computer could match given patches of terrain to them, i.e. how well a machine could automatically locate its position given an image which had been so decompressed. The flowchart on page 3 illustrates the conceptual design of the experiment.

The intent of the experiment was to obtain an indication of whether position location could be successfully accomplished by reconstructing a limited region from a compressed digital map information database and comparing it with a "window" of observed

information. In the experiment, a 512 x 512 pel image plays the role of the limited region of the digital map, and 64 x 64 pel test subimages play the role of the observed information.

```
┌────────────────────┐                            ┌──────────────┐
│   512 x 512 pel    │    Aware, Inc.             │  Compressed  │
│     Original       │  Image Compression         │   Overhead   │
│  Overhead Terrain  │ ─────────────────────────► │ Terrain Image│
│      Image         │                            │     File     │
└────────────────────┘                            └──────────────┘
```

Approximate
Compression Ratios
Employed in the
Experiment:
1:1
14:1
25:1
35:1
58:1
85:1
115:1

Aware, Inc.
Image Reconstruction

```
┌────────────────────────┐
│ Overhead Terrain image │
│   reconstructed  from  │
│  Compressed Image File │
└────────────────────────┘
```

Standard
Correlation
and Laplacian
Matching
Algorithms

```
┌────────────────────┐
│    64  X 64 pel    │
│      Terrain       │
│     Test Patch     │
│  (Uncompressed)    │
└────────────────────┘
```

Comparison
of
Position Location Accuracy
for
Reconstructed Compressed Terrain Image
vs.
Original Uncompressed Terrain Image

Test Image Windows

AQVIR1 W1
AQVIR1 W2
AQVIR1 W3
AQVIR1 W4
AQVIR 2 W1
AQVIR 2 W2
AQVIR 2 W3
AQVIR 2 W4

Flow Chart
for the
Position Location Experiment

Test patches were located and registered by maximizing the Laplacian of the surface defined by the correlation between the test patch and the reconstituted image. This process resulted in selection of the correct general location in every case for compression ratios through 117:1. No experiments were run at greater compression. To evaluate the amount of error introduced through the use of a decompressed image, sub-pixel registration was carried out, and the error was found to be less than one pel for all compression ratios.

Experiments were also run determining position location for the same data by maximizing the smoothed Laplacian and by maximizing the correlation function as criteria. The smoothed Laplacian worked in all cases. Use of the correlation function alone as a matching criterion led to correct position identification at all compressions for five cases. It failed for reference image compressions of 36:1 and greater for test patch AQVIR2W2, for compressions of 117:1 for test patch AQVIR2W3, and for compressions of 60:1 and greater for AQVIR2W4.

These results suggest that compression using wavelets preserves more than enough information to register position correctly using position reconstituted reference images that have been compressed in excess of 100:1. The results also illustrate how much important information is bound up in the Laplacian and other measures of the rate of change of the gradient, and how well this information is preserved by wavelet compression methods. A variety of analyses made to develop an indication of the relationship between information loss and compression are reported in the text.

Aware conducted two preliminary position location experiments in addition to those required to satisfy the contract. First, for AQVIR1, test patch 2 was corrupted with Gaussian random noise of average energy ranging from 0 to 24.9% that of the total image energy. In this experiment the test patch was still located with 100% accuracy at all compressions through 117:1 until the noise exceeded 20%.

Second, for AQVIR1, the data in test patch 1 were scaled from 80% to 120% of correct size; this test is equivalent to attempting to locate position with a difference in altitude of the same amount. The Laplacian correctly located the test patch within 2 pels at all compressions for scaling from -5% through +10%; within 3 to 4 pels for a scaling of +20%; within 5-6 pels for a scaling of -10%; and within 10 to 11 pels for a scaling of -20% to +20%. For all of the scalings but -10% and -20%, using maximization of the Laplacian as an initial step and then using sub-pixel registration based on maximization of the correlation surface will refine the result so that the error is *usually* less than 1 pel and *always* less than 2 pels.

A complete investigation of potential applications of these methods to position location problems posed as realistically as possible seems more than justified by these results. Such an investigation should include at least the following:

Other research conducted by Aware, inc. *suggests the feasibility of position finding by hierarchical correlation that by comparing observations with reference image data in the compressed space rather than the decompressed space normally used.* This offers the potential for using smaller amounts of high speed memory because it would not be necessary to decompress map sub-regions, and it should reduce the computational workload for the comparison algorithm.

(1)   a more realistic definition of the fitting problem;

(2)   a large and realistic set of reference images, test patch distortions, and noise parameters;

(3)   examination of rotational independence;

(4)   examination of the possibility of using hierarchical correlation to increase robustness and speed calculations; and

(5)   examination of ease of retrofit into existing systems.

Locating man-made objects in clutter

Identifying physical objects in images is a less well-formed problem than that of compressing images for use in determining position. Because it is less well-formed, the problem is paradoxically both more difficult to solve and easier to make modest progress toward solution. As an intermediate step, researchers seek reliable and cost-effective means to automate the process partially by cueing a human analyst.

One way in which natural objects differ from man-made objects is that the former often exhibit self-similar structures at various spatial scales whereas the latter normally exhibit symmetries at one scale rather than self-similarity at many scales. Trees, bodies of water, shorelines, and beaches are examples of physical objects whose constituents are similar throughout a variety of scales. Trucks, buildings, and roads are typical of structures that exhibit features at one or a few scales and are not self-similar. This difference between the structure of natural and manufactured objects is maintained in images of them for imaging wavelengths that are small compared with the size of the features of interest. These observations suggest that the wavelet representation might be effective in distinguishing between natural and manufactured structures in images because

the wavelet basis functions are functions whose graphs are related by translation and by similarity, where the scaling factor is 2. Hence, Aware hypothesized that the coefficients of the wavelet expansion of a function that exhibits self-similarity will exhibit a corresponding self-similarity if the natural self-similarity inherent in the wavelet basis is comparable to the self-similarity inherent in the function or image. Thus, subtraction of the self-similar parts of the wavelet expansion could be expected to highlight manufactured objects.

The objective of this part of the contract was to conduct an initial test of this hypothesis. For this test it was decided to employ representatives of several of the simplest wavelet bases in order to provide a foundation for possible more extensive future investigations.

Atlantic Aerospace provided two test images for this experiment, one an aerial photograph and the other a simulated synthetic aperture radar (SAR) image. Aware processed the images by computing their wavelet expansion, and selecting from their expansions those terms whose scale, i.e. support, was approximately *one-half*, *equal to*, and *twice* the size of the object(s) being sought. These scales could be expected to contain most of the image energy corresponding to the manufactured objects of interest, and they would also be expected to contain proportionally less of the image energy of the self-similar natural objects in the surrounding environment which, due to their self-similarity, would be expected to exhibit diffuse reflection over a broad range of illumination conditions.

According to the hypothesis, manufactured objects should display a greater difference from the mean image intensity than the self-similar natural background. Hence, partition of the pseudo-image pels into three categories appeared to be a reasonable step. The pseudo-images were "trinarized" by causing pels with values in approximately the middle 20% of the distribution — corresponding to the self-similar natural background — to be gray, those with smaller values — corresponding, according to hypothesis, to shadowed non self-similar structures — to be black, and those with greater values — corresponding to illuminated non self-similar structures — to be white.

Part I: Position Location using Stored Image Information

I.1 The Problem

The overall problem is that of determining location by comparing observations to a digitally stored representation of an aerial photograph reference image. Systems designers seek to minimize the amount of reference image information that must be stored to permit reliable, accurate location. These experiments examined how well test patches could be located against reference images decompressed using wavelet-based methods developed by Aware, Inc.

I.2 Description of the Demonstration

I.2.1 The Test Data.

Atlantic Aerospace Electronics Corporation provided Aware two digitized aerial photographic images to be used as reference images for this task. The images are designated 'AQVIR1' and 'AQVIR2'. Each is an 8 bits per pel gray scale image of 512x512 pels. Exhibits I-1 and I-2 are reproductions of the two images.

Atlantic Aerospace provided eight test patches, the position of which was to be determined by comparison with the compressed and reconstituted ("filtered") reference image. Each test patch is 64 x 64 pels. All four of the test patches for AQVIR1 and three for AQVIR2 were windows drawn from the master images. These test patches are designated

AQVIR1 W1, AQVIR1 W2, AQVIR1 W3, AQVIR1 W4
and
AQVIR2 W1, AQVIR2 W2, AQVIR2 W3.

Exhibits I-3 and I-4 show the location of the test patch windows in AQVIR1 and AQVIR2. In addition, a test patch designated AQVIR2 W4 was provided. AQVIR2 W4 was extracted from an image designated AQVIR5 that covers approximately the same region as AQVIR2 W3 but differs from it in that the time of day and angle of acquisition were slightly different. AQVIR2 W3 and AQVIR2 W4 are stereo image pairs. Since AQVIR2 W4 was not extracted from AQVIR2, it offers actual data that differs from the stored map both in viewing angle and lighting conditions and therefore represents a more realistic experimental test than the other windows provide. Exhibit I-5 shows test patch AQVIR2 W4 and also test patch AQVIR2 W3 for comparison.

Exhibit I-1 Image of AQVIR1

Exhibit I-2   Image of AQVIR2

Exhibit I-3   Image of AQVIR1 Showing Location of Test Patches

Exhibit I-4 Image of AQVIR2 Showing Location of Test Patches

Exhibit I-5 Image of Test Patch AQVIR2-4 with Image of Test Patch AQVIR2-3 Provided for Comparison

I.2.2   Procedure.

Aware used one of its wavelet-based image compression algorithms to compress the AQVIR1 and AQVIR2 test images . The compression ratios[1] used are listed in below.

Compression Ratios for AQVIR Image Tests

AQVIR1	1:1	14.2:1	23.5:1	33.7:1	58.1:1	90.1:1	116.7:1
AQVIR2	1:1	15.3:1	26.6:1	36.0:1	59.7:1	84.5:1	117.3:1

Compression ratio as used in this report is defined as the ratio of the number of bytes of computer memory required to store the compressed image data file, including all image-specific header information that is required by the decompression algorithm, to the number of bytes of computer memory required to store the original image data file. Although the *algorithms* used in Aware's image compression methods are independent of the image, they yield a *compression ratio* that has a small dependence on properties of the particular image. Thus it is not possible to precisely specify in advance the compression ratio that will result from given initial parameter settings for the compression algorithm.

This measure of compression ratio is conservative in the sense that it understates the theoretical compression ratio by taking into account the actual memory requirement and all other overhead for image storage in a real computing environment. The degree of conservatism is machine dependent to the extent that the machine word length may not be optimal for storing the compressed version of a given image.

Each compressed image data file was then decompressed, i.e. an image approximating the original image was constructed from the compressed image data file. We refer to this as the "decompressed" or "reconstituted" image. The AQVIR*n* decompressed image reconstructed from the data file corresponding to compression ratio R is designated AQVIR*n*[*R*]. For example, "AQVIR1[58]" designates the decompressed image of AQVIR1 reconstructed from the 58.1-to-1 compressed image data file.

The surfaces defined by the *correlation* between the reference image (uncompressed and decompressed for all compression ratios) and each test patch were then calculated, as well as the *thresholded Laplacian*, i.e. the Laplacian for each peak of the correlation surface whose value is greater than 20% of the "running maximum" of the correlation, and a *smoothed Laplacian*. The location of the test patch was then estimated in three different ways, by selecting the point at which one (or more) of the following occur: (a) the correlation assumes its maximum value, (b) the thresholded Laplacian assumes its maximum value, (c) the smoothed Laplacian assumes its maximum value. *Sub-pixel*

---

[1] The original AQVIR images do not use the full 0-255 range of intensity variation. This has the effect of reducing the compresion ratios quoted in the table by the factor 7.8/8.0 = 0.975 .

registration was then calculated for each selected position.

To analyze the results of the experiments, a set of 6 graphs was prepared for each image-test patch combination. To help visualize the correlation surface, radial diagrams were generated for all compression ratios showing the maximum value of the correlation surface at each radius and the value of the maximum of the Laplacian at each sidelobe. To better understand the effects of compression, graphs of peak-to-sidelobe ratio by compression were prepared both for the correlation function and its Laplacian. To place the effects of compression in information terms, graphs showing the uncompressed signal-to-noise ratio and the decline in signal-to-noise ratio with compression were also prepared for both the correlation and its Laplacian. Finally a few summary tables showing results for all images and test patches were prepared and are included with the analysis. All of the data supporting each graph and additional experimental data are included as an appendix.

I.2.3   Details of the Calculations

(i)   **Decompressed Image Intensity Function.**   Let the full image reconstituted from the file of one compressed by compression ratio CR be described by the image gray level function $I_{CR}(m,n)$ where $0 \leq m, n \leq 511$. The original uncompressed image has compression ratio $CR = 1$, so $I_1(m,n)$ is the image gray level intensity value for the pel labelled (m,n)

(ii)   **Correlation.**   A window   W   is represented by its gray level function $W(m,n)$ where $0 \leq m, n \leq 63$ for the 64 x 64 pel window. Only those image values that lie within the window   W   will contribute to the correlation. If the window is placed with its lower left corner (corresponding to window coordinates (0,0)) at the pel labelled (x,y), then the value of the correlation function is

$$\text{Cor}_{CR}(x,y) := \Sigma \; \{ W(m,n) - <W> \}\{ \; I_{CR}(x+m,y+n) \; - <I> \; \}$$

where the sum runs over all pairs (m,n) , $I_{CR}(x+m,y+n) = 0$ if $(x+m,y+n)$ lies outside the image, $<I>$ denotes the average vale of the intensity, and $<W>$ denotes the average value of the intensity of the window image. The quantities

$$\{W(m,n) - <W> \} \quad \text{and} \quad \{ \; I_{CR}(x+m,y+n) \; - <I> \; \}$$

represent pel intensity values that have been normalized to vary about zero. Note that cancellation of terms can occur for anticorrelated data.

The best match between the window and a portion of the image corresponds to that location (x,y) for which the correlation $\text{Cor}_{CR}(x,y)$ is maximal. If we think of the function (x,y) —> $\text{Cor}_{CR}(x,y)$ as a surface, then a likely position of the window corresponds to the location of the tallest peak of the surface. *The location of the maximum of $\text{Cor}_{CR}(x,y)$ is one measure of window location that was used in this demonstration.*

The array $W(m,n) - <W>$ can be thought of as a vector W, and the array $I_{CR}(x+m,y+n) - <I>$ can be thought of as a vector $I(x,y)$ indexed by the pair (x,y). Then the correlation is the inner product

$$\text{Cor}_{CR}(x,y) = W \cdot I(x,y) = \| W \| \| I(x,y) \| \cos \theta \; ;$$

hence the normalized correlation

$$\text{Cor}_{CR}(x,y) / \| W \| \| I(x,y) \| = W \cdot I(x,y) / \| W \| \| I(x,y) \|$$

lies between -1 and +1 inclusive, so

$$- \| W \| \| I(x,y) \| \le \text{Cor}_{CR}(x,y) \le \| W \| \| I(x,y) \| \; .$$

The normalized correlation may provide a more stable estimator than the un-normalized correlation $\text{Cor}_{CR}(x,y)$, although the results reported herein indicate that normalization may not be required for robust position location.

(iii) **Laplacian.** A surface is the graph of a function of two variables. At an isolated peak of a surface the vector of first partial derivatives of the function is zero and the matrix of second partial derivatives describes the rate with which the surface drops off from the maximum in different directions. A Taylor series approximation to the function in a neighborhood of the peak will consist of a constant plus a quadratic form in x and y whose coefficients are the second partial derivatives. The quadratic terms describe the intersection of a plane parallel to the x - y plane that intersects the surface near the peak point. The intersection curve will be an ellipse. The axes of the ellipse are given by the eigenvalues of the matrix of second partial derivatives. The eigenvalues will be equal if and only if the intersection curve is a circle, and in this case the matrix of second partial derivatives corresponds to the Laplacian operator,

$$\Delta f := \partial^2 f / \partial x^2 + \partial^2 f / \partial y^2$$

The Laplacian, i.e. the absolute numerical value of the rotation-invariant second order

differential operator derived from the matrix of second derivatives of the correlation function $Cor_{CR}(x,y)$ evaluated at the nominal peak, can be useful in locating a match between two surfaces. The Laplacian is approximated by the 5 point difference scheme

$$| \Delta f(m,n) | = | f(m+1,n) + f(m-1,n) + f(m,n+1) + f(m,n-1) - 4f(m,n) | .$$

To avoid being led astray by very large Laplacians associated with minute spikes of the correlation surface, it is useful to use a thresholding criterion so that a local maximum of the Laplacian is ignored unless the spike involved has a height greater than either some pre-assigned value or estimator based on some characteristic of the data. In this experiment, the thresholding algorithm ignored local maxima of the Laplacian if the local maximum of the correlation surface at which the Laplacian was evaluated was less than 20% of the global maximum of the correlation. Thresholding was used as part of the algorithm to locate the position of the test patches by seeking the maximum Laplacian. This

Peak Correlation

20% of Peak Correlation

is illustrated in the figure, where only those peaks that rise above the "20%" dashed line are candidates for the maximum Laplacian calculation.

(iv) **Smoothed Laplacian.** The Laplacian was estimated directly from the numerical data and also by employing a Fourier smoothing operator followed by interpolation. It was hypothesized that such a smoothed Laplacian would be less affected by noise. The smoothing operator selected was based on the Shannon-Whittaker interpolation formula, which expresses a function whose Fourier transform has bounded support in terms of the values of the function sampled on a lattice. For a function of one variable the formula is

$$f(x) = \Sigma f(n) \sin \pi(x-n) / \pi(x-n)$$

where f is sampled on the lattice of integers and the Fourier transform of f is zero outside the interval [-1/2,1/2] . The correlation function can be smoothed by assuming that the pixel values are sampled values on the lattice of integer coordinate pairs (m,n) and that the

Fourier transform of the sampled correlation function is zero outside the product interval [-1/2,1/2] x [-1/2,1/2] . With these assumptions, the interpolation formula is

$$f(x,y) = \sum f(m,n) \left(\sin \pi(x-m) \sin \pi(x-n) \right) / (\pi^2(x-m)(x-n))$$

With of this formula the value of the function f can be estimated for any point (x,y) .

The interpolation formula was implemented by calculating the Discrete Fourier Transform of the correlation surface for the 8 x 8 neighborhood of the peak pixel, then embedding the four quadrants in a 64 x 64 matrix, and computing the inverse Fourier transform. Since the effective sampling rate is 64^2 samples per 8 x 8 neighborhood, this method provides interpolation that is accurate to within 1/8 pixel.

**Correlation Surface in Neighborhood
of
Maximum of Laplacian**

(v) **Sub-Pixel Registration.** If the correlation matrix is thought of as a collection of sampled values of a continuous function of two variables, then the peak of the correlation surface will lie within some pixel but its coordinates need not, in general, coincide with the integer coordinates on which the stored image was recorded. The precise location of the peak can be approximated by interpolation of the surface to find its maximum. The interpolated coordinates of the tallest peak provide sub-pixel registration of the window in the image. The same spectral interpolation formula was used for sub-pixel registration as for smoothing the Laplacian. Sub-pixel registration was performed for peaks located by maximizing the Laplacian and for those located by maximizing the correlation function using the same criterion function.

(vi) **Peak-to-Sidelobe Ratio.** The *peak-to-sidelobe ratio* is the ratio of the height of the estimated maximum of the correlation surface to that of the next largest sidelobe, key information in making the registration decision. The *sidelobe-to-peak ratio* := 1/(peak to sidelobe ratio).

I.3 Results and Analysis

Exhibit I-6 summarizes the results of the experiment. A test case is said to have "failed" if the window position was not correctly determined within an accuracy of one pel. For failed test cases the lowest compression ratio at which failure occurred is given. For instance, maximizing correlation failed to correctly locate window AQVIR2 W2 to within 1 pel at compression ratio 36:1.

As shown by Exhibit I-6, maximization of the Laplacian (both unsmoothed and smoothed) located the position of the test patch within one pel or less of the correct position in every case for all compression ratios tested (0-117).[2] In three of the eight test cases, after some degree of compression, maximization of the correlation function led to selection of an incorrect position.

AQVIR2 W3 and AQVIR2 W4 are an image pair, with AQVIR2 W4 drawn from a reference image which differs from AQVIR1 in angle of acquisition and illumination. It is interesting to note that both Laplacian methods succeeded in locating AQVIR2 W4, and that the maximum correlation method failed at the relatively high compression ratio of 59:1.

Exhibit I-6
Compression Ratios at Failure by
Test Case and
Method of Positioning

Test Case	Laplacian	Smoothed Laplacian	Correlation
AQVIR1 W1	None	None	None
AQVIR1 W2	None	None	None
AQVIR1 W3	None	None	None
AQVIR1 W4	None	None	None
AQVIR2 W1	None	None	None
AQVIR2 W2	None	None	36:1
AQVIR2 W3	None	None	117.3
AQVIR2 W4	None	None	59.

The Laplacian measures retained the ability to accurately locate position within images decompressed using Aware's wavelet-based image compression method at high compression ratios. These results suggest that there may be a potentially significant savings in reference map digital storage requirements without compromising reliability for

---

[2] The only three cases where the error reached one pel were AQVIR1W1(116), AQVIR1W4(116), and AQVIR2W3(84), all for the unsmoothed Laplacian. For AQVIR2W3, compression ratios greater than 84 yielded registrations accurate to within less than one pel.

realistic digital map-based position location systems. Validation of Aware's method would require additional study, including a statistical investigation based on a more realistic problem specification.

The reason for these excellent results can be seen by examining Exhibits I-7 through I-12. All of these exhibits are graphs with *compression* as the abscissa. On each graph are eight lines representing the behavior of a specific reference image test patch combination. The labels on these lines are of the form $AmWn$, where An is reference image $AQVIRm$ and Wn is test patch n for reference image Am. The exhibits are sequentially arranged in groups of two; the upper, odd-numbered exhibits show how the correlation surface behaved and the lower, even-numbered exhibits show how the Laplacian behaved. The first pair of exhibits shows the peak values of the correlation and Laplacians; the second pair of exhibits shows the ratio of the peak values to the values at the next largest side lobes; and the third pair shows the SNR of the correlation function and of the Laplacian.

Exhibit I-7 Correlation at Peak vs Compression

Y-axis: Normalized Correlation at Correct Location (0 to 1)
X-axis: Compression Ratio (20 to 100)

Curve labels (right side, top to bottom): A1w2, A1w1, A1w3, A2w1, A2w3, A2w1, A1w4, A2w2

Exhibit I-8

Laplacian at Peak vs Compression

Y-axis: Max Laplacian Value for Thresholded Correlation (x10^6) (0 to 7)
X-axis: Compression Ratio (20 to 100)

Curve labels (left side, top to bottom): A1w4, A1w3 A2w1, A1w1 A2w3 A2w2, A1w3, A2w4

Exhibit I-9

Ratio of Peak Correlation to that of Largest Side Lobe by Compression

Compression Ratio

Exhibit I-10

Ratio of Peak Laplacian to that of Largest Side Lobe by Compression

Compression Ratio

Exhibit I-11

SNR in Correlation by Compression (Largest SideLobe)

Compression Ratio

Exhibit I-12

SNR in Laplacian by Compression (Largest SideLobe)

Compression Ratio

Six of the test samples were numbered in sequential order of difficulty; namely AQVIR1W1–4, then AQVIR2W1–2. As previously noted, AQVIR2W3 and AQVIR2W4 are a pair of test patches selected to test the robustness with which location could be determined when the reference image and the test patch differed in a likely way, namely that the images had been acquired at different times and from different angles. These test patches appear to have a base level of difficulty between AQVIR2W1 and AQVIR2W2.

More importantly, in Exhibits I-11 and I-12, one can see a relatively uniform decrease in the SNR of both the correlation function and the Laplacian. The SNR of correlation decreased at a mean rate of approximately 1.5db per 100 compression units with a range in the rate of decrease from 1.0db per 100 compression units to 2.5db per 100 compression units and a very strong tendency that the smaller the uncompressed peak-to-sidelobe ratio, the faster the rate of decline. Given that the correlations between AQVIR2 and AQVIR2W2–4 all had uncompressed peak-to-sidelobe ratios less than 2, it was not possible at high compression to correctly locate their position by maximizing correlation.

The SNR of the Laplacian decreased consistently and nearly uniformly across the test cases at a rate of approximately 4.2db per 100 compression units. Given that the uncompressed peak-to-sidelobe ratio of the Laplacians for all the test cases fell between 7db and 8.5db, maximizing the Laplacian should be more than sufficient to identify the correct location for compressions past 120. In hindsight, we should not be surprised that Aware's compression method preserved Laplacian information so well under compression. The method uses wavelets of bounded support as a basis for the compression and as a result retains transients better than other methods, i.e. it smooths the derivative or Laplacian less.

As mentioned above, A2W4 is a test patch based on an aerial photograph taken at a different time of day and from a different angle of the same area as test patch A2W3, which was drawn from the reference image. Using either the Laplacian or smoothed Laplacian yielded A2W4's correct location at all compressions. As shown in Exhibit I-11, SNR of the correlation of A2W4 with A2 averaged about 0.5db *more* than that for A2W3. Thus, using correlation alone to locate A2W4 worked better than it did for A2W3, failing at compression 79 rather than 32. This may be a statistical anomaly due to the small size of the test data sample.

These results are particularly important, because the A2W4 test involves a realistic kind of variation between test patch and reference image. That A2W4's variant test patch should have been easier to locate than A2W3's perfect match appears counter-intuitive. However, if one assumes that the variation from A2W3 to A2W4 had the effect of raising the noise floor, then the side lobe would be swallowed up before the main lobe and the peak-to-sidelobe ratio would improve even though the absolute strength of the main signal declined.

I.4   Additional Experiments

Practical application requires that a compression method for reference images permit robust location when the test sample has been corrupted by noise, scale change, and rotation. Although not required to do so by the contract, Aware conducted preliminary inquiries into the effects of *noise* and *scale change*. No experiments were conducted with respect to rotation.

I.4.1   Noise

The preliminary study of effects of noise on ability to locate was conducted with decompressed versions of reference image AQVIR1 and test patch AQVIR1 W2. Noise was added only to A1W2, and the stored decompressed reference image was kept noise-free. Only one sample of noise was used per noise level. Although the use of but one noise sample per test limits the statistical validity of the experiment, the results are sufficiently promising to warrant a more detailed statistical study.

The value of each pixel in the noisy test patch W* was computed thus

$$W*(m,n) := A1W2(m,n) + Noise(m,n)$$

where Noise(m,n) is computed thus

$$Noise(m,n) := \sqrt{R}\sqrt{P}\ Rand(m,n)$$

where Rand(m,n) is a matrix of Gaussian random numbers with mean 0 and variance 1, P denotes the power of image to which the noise is added, R is the ratio of the noise power to the image power, and image power is calculated thus

$$P := \Sigma_{m,n}[A1W2(m,n)]^2$$

The noise was varied from 1% to approximately 24%, and an attempt was made to locate the noisy test patch for each of the compression ratios of the reference image used in the main study. "Percent noise" is defined by:

Percent Noise: = 100 * Power(Noise)/ [Power(Image) + Power(Noise)]

For all compressions and for all noise levels through 23.6%, all three methods yielded the correct location. At noise level 24.9%, it was not possible to arrive at the

correct location. This result may reflect the specific characteristics of the pseudo-random number generator, as a quick trial at a noise level in excess of 30% once again yielded the correct location.

To see what bias is introduced by the combination of noise and compression, sub-pixel registration error was calculated for all 7 compressions and 11 noise levels when location was determined by seeking the maximum of the Laplacian of thresholded correlation. Exhibit I-13 displays the result. Sub-pixel registration error varied between zero and 1.152 pels for noise levels not greater than 23.6%. Clearly no significant bias is introduced.

Exhibit I-13
Sub-Pixel Registration Error in Pels of
Test Patch A1W2 (with Varying Amounts of Noise) with
Reference Image A1 (at Various Compression Ratios)

Noise Level	Compression Ratios						
	1:1	14:1	23:1	33:1	58:1	70:1	120:1
1%	0.0	.800	.976	1.061	1.152	1.152	.901
4.6%	0.0	.800	.976	1.061	1.152	1.152	.901
8.3%	0.0	.800	.976	1.061	1.152	1.152	.901
11.5%	0.0	.800	.800	1.061	1.152	1.152	.901
14.2%	0.0	.800	.800	1.061	1.152	1.152	.901
16.6%	0.0	.800	.800	1.061	1.061	1.152	.901
18.7%	0.0	.800	.800	1.061	1.061	1.152	.901
20.5%	0.0	.800	.800	1.061	1.061	1.152	.901
22.1%	0.0	.800	.800	1.061	1.061	1.152	.901
23.6%	0.0	.800	.800	1.061	1.061	1.152	.901
24.9%	0.0	*	*	*	*	*	*

* Original registration was more than 100 pels in error; sub-pixel registration did not help.

A sense of how robustly location can be determined by using these means of compression is suggested by perusing Exhibits I-14-a through -g which show the Laplacian for all noise levels at each compression.(*Note*: the scale differs from graph to graph.)

Exhibit I-14-a

Laplacian of Correlation of A1 (1:1) with A1W2 (various noise levels) vs. Distance from Maximum Peak

Exhibit I-14-b

Laplacian of Correlation of A1 (14:1) with A1W2 (various noise levels) vs. Distance from Maximum Peak

Exhibit I-14-c

Laplacian of Correlation of A1 (23:1) with A1W2 (various noise levels) vs. Distance from Maximum Peak

Curves for % noise levels 1.0 4.6 8.3 11.5 14.2 16.6 18.7 20.5 22.5 23.6

Radial Distance from Max Peak

Exhibit I-14-d

Laplacian of Correlation of A1 (33:1) with A1W2 (various noise levels) vs. Distance from Maximum Peak

Curves for % noise levels 1.0 4.6 8.3 11.5 14.2 16.6 18.7 20.5 22.5 23.6

Radial Distance from Max Peak

Exhibit I-14-c

Laplacian of Correlation of A1 (58:1) with A1W2 (various noise levels) vs. Distance from Maximum Peak

Radial Distance from Max Peak

Exhibit I-14-f

Laplacian of Correlation of A1 (90:1) with A1W2 (various noise levels) vs. Distance from Maximum Peak


Radial Distance from Max Peak

Exhibit I-14-g

Laplacian of Correlation of A1 (117:1) with A1W2 (various noise levels) vs. Distance from Maximum Peak

## I.4.2   Scaling

A preliminary study of effects on position determination of a difference in scale between the test patch and the decompressed reference image was conducted using reference image AQVIR1 and test patch AQVIR1 W1. Holding the decompressed reference image scale constant, the test patch was stretched using Fourier interpolation at fourteen different scale ratios ranging from 0.8 to 1.2.

Again, as shown in Exhibit I-15 the Laplacian correctly located the test patch within 2 pels at all compressions for scale changes from -5% through +10%; within 3 to 4 pels for a scaling of +20%; within 5-6 pels for a scaling of -10%; and 10 to 11 pels for a scaling of -20% to +20%.

Exhibit I-15
Location Error in Pels for
Test Patch A1W1 (Scaled from 0.8 to 1.2)
Located by Maximization of Laplacian of Correlation with
Reference Image A1 (Variously Compressed)

	Compression Ratio					
Scaling	1:1	23:1	33:1	58:1	90:1	120:1
.80	10.30	10.30	10.44	10.44	9.43	9.90
.90	5.39	5.39	5.39	5.39	5.00	5.09
.95	2.24	2.24	2.24	2.24	2.24	2.24
.97	1.00	1.41	1.41	1.41	1.41	1.00
.98	1.41	1.41	1.41	1.41	1.00	1.00
.99	0.00	0.00	0.00	0.00	0.00	0.00
1.00	0.00	0.00	0.00	0.00	0.00	1.00
1.01	0.00	0.00	0.00	0.00	0.00	1.00
1.02	0.00	0.00	0.00	0.00	0.00	1.00
1.03	0.00	0.00	0.00	0.00	0.00	1.00
1.05	0.00	0.00	0.00	1.00	0.00	1.00
1.10	1.41	1.41	1.41	1.41	1.00	1.00
1.20	2.83	2.83	2.83	2.83	3.61	4.24

The large errors in Exhibit I-15 for -20% and -10% scale reduction of the test patch correspond to selection of the wrong peak in the correlation function, rather than to local interpolation errors. As shown in Exhibit I-16, for all of the scalings but -20% and -10%, using maximization of the Laplacian as an initial step and then using sub-pixel registration based upon maximization of the correlation surface will refine the result so that the error is usually less than 1 pel and always less than 2 pels. In brief, if the noise level is not so great that it induces selection of the wrong peak, then interpolation will find the correct

location to within 2 pels.

Exhibit I-16
Sub-Pixel Registration Error in Pels of
Test Patch A1W1 (scaled from 0.8 to 1.2) with
Reference Image A1 (compressed from 1:1 to 120:1)

Compression Ratios

Scaling	1:1	14:1	23:1	33:1	58:1	90:1	120:1
.80	1.944	1.944	1.944	1.944	1.944	1.944	1.944
.90	1.944	1.944	1.944	1.944	1.944	1.944	1.944
.95	.451	1.944	1.944	1.944	1.944	1.944	1.944
.97	.707	.707	.625	.625	1.944	1.944	1.425
.98	.707	.707	.625	.625	1.944	1.944	.707
.99	.451	.707	.707	.707	.707	.707	.707
1.00	0.000	.515	.707	.707	.707	1.008	1.061
1.01	1.768	.707	.707	.707	.707	.707	.976
1.02	0.000	.707	.707	.707	.707	.951	.976
1.03	.125	.707	.707	.707	.707	.910	.901
1.05	.673	.707	.839	.951	.910	.910	.901
1.10	1.068	1.068	.910	.910	1.944	.910	.910
1.20	.625	1.944	1.944	1.944	1.944	1.944	1.068

An indication of how scaling affects the ability to locate test patches relative to images can gained by perusing Exhibits I-17, I-18, and I-19, which for all compression ratios display the value of the unnormalized maximum correlation versus scaling factor, the value of the maximum Laplacian versus scaling factor and the value of the maximum smoothed Laplacian versus scaling factor, respectively.

Exhibit I-17

AQVIR1 W1: Maximum Correlation Value vs Scaling Factor
Curves for compression ratios 1:1, 14:1, 23:1, 33:1, 58:1, 90:1, and 117:1

Window Scaling Factor

Exhibit I-18

AQVIR1 Window 1: Laplacian vs Window Scaling Factor

Exhibit I-19

AQVIR1 Window 1: Smoothed Laplacian vs Window Scaling Factor

I.5   Conclusions and Recommendations

The consistent predictability and slow decline with compression of the peak-to-sidelobe ratio of the Laplacian and of the Coefficient of Correlation is noteworthy. The results were encouraging even in the presence of distortion and noise, as seen in the experiment with AQVIR2 test patch W4, and the preliminary noise and scaling experiments exceeded expectations. That a test patch corrupted by more than 20% noise could be correctly located relative to a reference image reconstituted from a file compressed 120:1 indicates that Aware's wavelet-based compression methods show potential for practical applications.

In addition, *other research conducted by Aware, Inc., suggests the feasibility of position finding by hierarchical correlation, that is by comparing observations with the reference image data in the compressed space rather than the decompressed space normally used.* This offers the potential for using smaller amounts of high speed memory because it would not be necessary to decompress map sub-regions, and it should reduce the computational workload for the comparison algorithm.

A complete investigation of potential applications of these methods to position location problems posed as realistically as possible seems more than justified by these results. Such an investigation should include at least the following:

(1)   a more realistic definition of the fitting problem;

(2)   a large and realistic set of reference images, test patch distortions, and noise parameters;

(3)   examination of rotational independence;

(4)   examination of the possibility of using hierarchical correlation to increase robustness and speed calculations; and

(5)   examination of ease of retrofit into existing systems.

Part II: Identifying Objects in Clutter

II.1 Statement of the Problem

Part II describes results of an initial study of the use of wavelets to locate manufactured objects in clutter. As presented to Aware, Inc., the experiment concerned computer-assisted search for vehicle-sized manufactured objects in overhead imagery.

Atlantic Aerospace provided two images for analysis, TANKS and HT1. TANKS is a daylight aerial photograph of desert which shows terrain features of varying size, one tank, one personnel carrier, and numerous vehicle tracks. HT1 is a simulated SAR image of a region with trees, a road, and tanks.
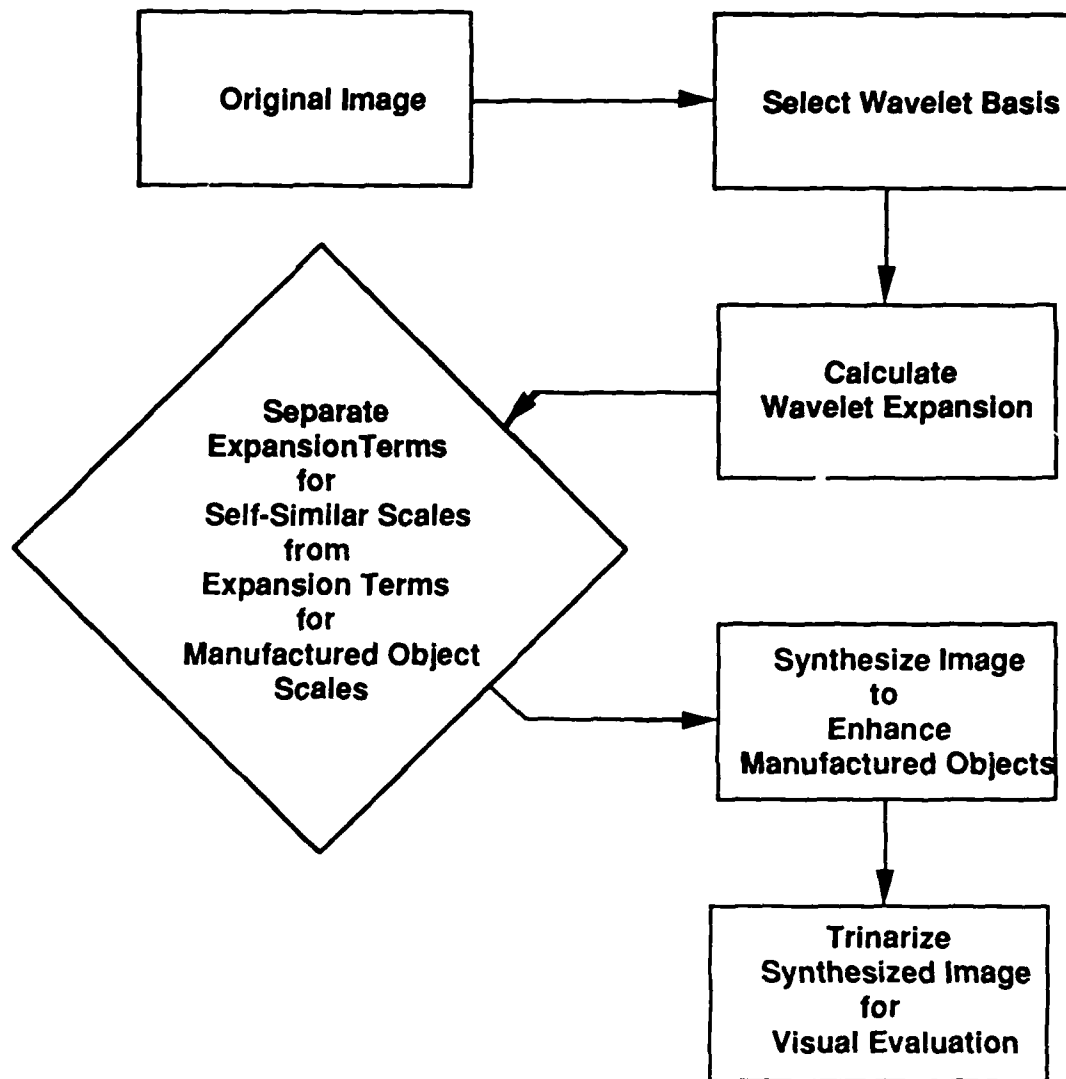
II.2 Description of the Experiment

Aware created experimental images from the original data by (1) computing the wavelet expansion of the image intensity function; (2) selecting the three scales in the wavelet expansion that were closest to the scale of the objects of interest (namely, one-half the scale, the same scale, and twice the scale); and (3) synthesizing images from the selected terms of the wavlet expansion. It was postulated that these synthesized images would exclude clutter due to objects in the image whose size is generally different from the size of the objects of interest, thereby reducing the complexity of finding them.

It was also postulated that by quantizing the gray scale levels of the synthetic images into a reduced number of levels, differences between manufactured and natural objects would be highlighted. Thus, the synthesized image gray levels were "trinarized" by quantizing them into three categories according to the following procedure: The mean of the pixel values was computed. The trinarization was then based upon two fractions, an "above" fraction, which determined a threshold value between the mean pixel value and the maximum pixel value, and a "below" fraction, which determined a threshold value between the minimum pixel value and the mean pixel value. These fractions determine three bins into which the pixels are then placed. The value assigned to pixels with values in the lowest tier is the minimum pixel value; in the middle tier, half the maximum value; and in the top tier, the maximum value. Some tuning by eye based on the actual distribution of pixel values can be helpful in selecting the "above" and "below" fractions.

A flow chart for the experiment is shown on the next page.

```
┌──────────────────┐              ┌──────────────────┐
│                  │              │                  │
│  Original Image  │─────────────▶│ Select Wavelet   │
│                  │              │      Basis       │
└──────────────────┘              └──────────────────┘
                                           │
                                           ▼
                                  ┌──────────────────┐
          ◇                       │                  │
    Separate          ◀───────────│   Calculate      │
  ExpansionTerms                  │ Wavelet Expansion│
      for                         │                  │
 Self-Similar Scales              └──────────────────┘
     from
 Expansion Terms
     for                          ┌──────────────────┐
 Manufactured Object              │ Synthesize Image │
     Scales         ─────────────▶│       to         │
          ◇                       │    Enhance       │
                                  │Manufactured Objects│
                                  └──────────────────┘
                                           │
                                           ▼
                                  ┌──────────────────┐
                                  │    Trinarize     │
                                  │ Synthesized Image│
                                  │       for        │
                                  │ Visual Evaluation│
                                  └──────────────────┘
```

Flow Chart
for
Objects in Clutter Experiment

The experimental images were prepared using four different types of wavelet bases designated "Haar", "Hat," "NN", and "D6". These wavelets are illustrated in Exhibits I-20 through I-23 respectively. They were selected because each is characteristic of a class of wavelet bases defined by a scaling coefficient recursion with 6 or fewer contiguous non-zero coefficients, and each of the four classes can be expected to highlight different categories of objects. The Haar scaling function and fundamental wavelet have the smallest possible support — 1 unit — and can be expected to be most sensitive to transient phenomena.

The length of the support of the scaling function and fundamental wavelet for each of the other three wavelet bases are the same, 6 units. These wavelet bases differ primarily in smoothness. The scaling function for D6 is differentiable and a series of its translates can exactly represent any polynomial of degree 2. The scaling function for Hat is continuous and a series of its translates can exactly represent a polynomial of degree one. Moreover, the Hat basis provides an expansion that "mimics" properties of the Laplacian of the Gaussian that has been used with some success for edge detection and representation. The basis functions of the wavelet basis designated "NN" have Fourier transforms that have an unusually large fraction of their energy at high frequencies, and provide an alternative method for characterizing edge-localized concentrations.

It was expected that application of D6 would lead to clutter reduction while preserving the general similarity of the synthetic and the original images. It was also expected that the employment of Hat might boost the relative energy in edges and, because the edges of manufactured objects tend to have greater linearity than those of natural objects, that the manufactured objects would be preferentially cued by the Hat-based synthetic image. Finally, NN was employed as a test to develop a better general understanding of the kinds of information that would be preserved by the synthetic expansion. Its utility for analysis will depend on analytic invariants that may be preserved or enhanced by this representation, and its effectiveness remains to be determined.

**Exhibit I-20: The Haar Wavelet**

Exhibit I-23: The D6 Wavelet

**Exhibit I-21: The Hat Wavelet**

Exhibit I-22: The NN Wavelet

II.3 Results and Analysis

For both of the original images and each of the four synthetic images, the results of the processing are shown on two photographs. The first photograph, called the *wavelet expansion photograph,* shows three levels of the wavelet expansion. The second photograph, called the *trinarization photograph,* shows the effect of trinarizing each of the decompressed images shown in the wavelet expansion photograph. Each photograph is divided into four quadrants numbered by the traditional numbering scheme as shown in the diagram below.

Quadrant II always displays the original image as provided to Aware, Inc. Quadrant III shows the wavelet expansion terms corresponding to wavelets of "wavelength" twice that of the sought objects. Quadrant IV shows the wavelet expansion terms corresponding to wavelets of twice the size and the same size. Quadrant I shows the wavelet expansion terms corresponding to wavelets of twice the size, the same size and half the size. Exhibits I-24 through I-31 are reproductions of these photographs.

Exhibit I-24: The TANKS Image Processed with the Haar Wavelet

   a. Effects of Compression

b. Effects of Trinarization

Exhibit I-25: The TANKS Image Processed with the Hat Wavelet

a. **Effects of Compression**

b. Effects of Trinarization

Exhibit I-26:   The TANKS Image Processed with the NN Wavelet

   a.   Effects of Compression

b. **Effects of Trinarization**

Exhibit I-27:   The TANKS Image Processed with the D6 Wavelet

a.   Effects of Compression

b. Effects of Trinarization

Exhibit I-28: The HT1 Image Processed with the Haar Wavelet

a. Effects of Compression

b. Effects of Trinarization

Exhibit I-29: The HT1 Image Processed with the Hat Wavelet

    a.   Effects of Compression

b. Effects of Trinarization

Exhibit I-30: The HT1 Image Processed with the NN Wavelet

a. Effects of Compression

b. Effects of Trinarization

Exhibit I-31: The HT1 Image Processed with the D6 Wavelet

   a. Effects of Compression

b. **Effects of Trinarization**

Overall the methods appeared to provide an enhancement to the conventional aerial photograph that would enable an analyst to concentrate attention on objects of the correct scale, particularly when the D6 wavelet was used. As shown by Exhibit I-27, the regions which were delimited using the trinarization technique quite accurately reflected the outlines of the objects present in the image. By comparing the highlighted images at the different scales, virtually all of the background clutter could be removed, leaving only objects of the desired scale. As shown by Exhibit I-25, the Hat wavelet produced less useful results. As shown by Exhibits I-24 and I-26, the Haar and NN wavelets were largely unsuccessful.

The success of the D6 wavelet basis can be attributed to its preservation of low frequency information *relative to its scale*. This suggests that other wavelet bases that have a similar property might exhibit equal or better performance. Further study of these alternatives would be worthwhile, with particular attention being paid to the trade-offs amongst the properties of good localization, minimal computational cost, and most effective cueing.

These simple techniques appear to have been largely unsuccessful when they were applied to SAR imagery. We believe that the wavelet-based methods worked for optical spectrum aerial photographs but not for SAR because information in optical spectrum aerial photographs tends to have important regional clusters of information that are *localized in scale* whereas the SAR process yields images that are constructed as loosely associated sets of edges whose aggregate relationships contain the important information. A raw SAR image does not contain information at coarse scales. As an edge *per se* has no distinct scale in the way that a region does, neither image expansion using wavelets of the sought after scale nor trinarization could of themselves be helpful.

A better approach for the interpretation of SAR imagery might be to enhance the edges present in the images and then use an artificial intelligence technique to group them by region growing. This procedure will produce image regions at various scales implied by the raw SAR data, within which the sought for manufactured objects are likely to appear. This processed image can be used as the input for a scale-selecting wavelet expansion which can be followed by trinarization or some comparable scale-based clutter removal process. We believe that this approach deserves study.

II.4 Conclusions and Recommendations for Future Work

Further investigation of these methods using Daubechies family wavelets could be fruitful for images of photographic type. If it were pursued, it might be helpful to combine them with an edge detector and a morphology analyzer.